

NASA Contractor Report 189660

ICASE

SEMIANNUAL REPORT

October 1, 1991 - March 31, 1992

Contract No. NAS1-18605
May 1992

Institute for Computer Applications in Science and Engineering
NASA Langley Research Center
Hampton, Virginia 23665-5225

Operated by the Universities Space Research Association



National Aeronautics and
Space Administration

Langley Research Center
Hampton, Virginia 23665-5225

N92-30459

Unclas

G3/59 0106605

(NASA-CP-100660) SEMI-ANNUAL FINAL
REPORT, 1 OCTOBER 1991 - 31 MARCH
1992 (ICASE) 49 p

CONTENTS

	Page
Introduction	ii
Research in Progress	1
Reports and Abstracts	26
ICASE Interim Reports	36
ICASE Colloquia	37
ICASE Staff	39

INTRODUCTION

The Institute for Computer Applications in Science and Engineering (ICASE) is operated at the Langley Research Center (LaRC) of NASA by the Universities Space Research Association (USRA) under a contract with the Center. USRA is a nonprofit consortium of major U. S. colleges and universities.

The Institute conducts unclassified basic research in applied mathematics, numerical analysis, and computer science in order to extend and improve problem-solving capabilities in science and engineering, particularly in the areas of aeronautics and space research.

ICASE has a small permanent staff. Research is conducted primarily by visiting scientists from universities and industry who have resident appointments for limited periods of time as well as by visiting and resident consultants. Members of NASA's research staff may also be residents at ICASE for limited periods.

The major categories of the current ICASE research program are:

- Numerical methods, with particular emphasis on the development and analysis of basic numerical algorithms;
- Control and parameter identification problems, with emphasis on effective numerical methods;
- Computational problems in engineering and the physical sciences, particularly fluid dynamics, aeroacoustics, and structural analysis;
- Computer systems and software for parallel computers.

ICASE reports are considered to be primarily preprints of manuscripts that have been submitted to appropriate research journals or that are to appear in conference proceedings. A list of these reports for the period October 1, 1991 through March 31, 1992 is given in the Reports and Abstracts section which follows a brief description of the research in progress.

¹Presently, ICASE is operated at NASA Langley Research Center, Hampton, VA, under the National Aeronautics and Space Administration, NASA Contract No. NAS1-18605. In the past, support has been provided by NASA Contract Nos. NAS1-18107, NAS1-17070, NAS1-17130, NAS1-15810, NAS1-16394, NAS1-14101, and NAS1-14472.

RESEARCH IN PROGRESS

APPLIED AND NUMERICAL MATHEMATICS

Numerical Analysis

Saul Abarbanel

The study of stable boundary conditions for compact high order scheme continued. Together with David Gottlieb and Mark Carpenter (Fluid Mechanics Division - Theoretical Flow Physics Branch, LaRC) the scalar formulation has been generalized to the system case.

The work (together with doctoral student Jeff Danowitz) on non-reflecting boundary conditions for compressible viscous flow past two-dimensional finite bodies has now yielded first practical results; working codes for subsonic airfoils have in the past been forced to incorporate the (mathematically) incorrect numbers of far down-stream boundary conditions. This situation has now been corrected.

David Gottlieb

I have continued to study the asymptotic stability of compact schemes. We are trying to find an energy norm for those schemes to assure stability also in the system case. The idea here is to change the equations in the neighborhood of the boundaries to assure stability in an energy norm. We have devised a method of formulating the equations in order to find the norm.

We started an effort in simulating interactions of shock waves and hydrogen bubbles. The techniques is spectral shock capturing. Some promising results indicate the validity of the method.

We also study the possibility of using Wavelets as basis functions for the numerical solutions of PDE's. We have identified some of those techniques as known finite difference schemes.

We have also continued the work on domain-decomposition spectral methods. We have found some ways of treating the interface boundary condition to enhance parallelism. In particular, a mixed explicit implicit scheme for parabolic equations looks very efficient.

John Van Rosendale

Work continues on multigrid algorithms which are robust, in the sense of achieving fast convergence on highly stretched grids, yet map well to highly parallel architectures. As is well

known, multigrid algorithms based on line- or plane-relaxation can achieve fast convergence independent of grid stretching and cell aspect ratios. Unfortunately such line- or plane-relaxation algorithms parallelize very poorly, especially in three dimensions, where only plane-relaxation suffices in general.

A better alternative, developed recently by the author and Naomi Naik of Columbia, is use of point relaxation together with multiple coarse grids formed by semi-coarsening the grid separately in each of the coordinate directions. This idea, closely related to the hyperbolic algorithm of Wim Mulder, can be used to create robust and effective elliptic solvers. In recent work, the authors have established the rapid convergence of several of these algorithms, and have demonstrated their effectiveness on a sequence of test problems. Our most recent result in this direction is a V-cycle proof for one version of this algorithm, in the more difficult case of variable coefficients and stretched grids.

In related work, Joe Dendy of Los Alamos National Laboratory is collaborating on the development of robust multigrid algorithms for problems having severe coefficient jumps. By combining the new algorithm with interpolation formulas developed previously for plane relaxation algorithms, we can now handle problems in which the coefficients vary by many orders of magnitude from cell to cell. In coming months we will implement a test code based on this algorithm on the CM 2 at Los Alamos and probably also on their CM 5, once the floating point coprocessors have been installed.

These new multigrid solvers use only point relaxation so implementation on highly parallel architectures is quite natural, although there are mapping issues to be resolved. Unfortunately, our most robust algorithm for problems with coefficient jumps still requires line-solves in the formulation of the interpolation operators for two dimensional problems. In three dimensions these line-solves become plane-solves, which are quite expensive, but no alternative to these line- and plane-solves has yet been discovered which yields comparable robustness.

A third collaboration on the same general class of algorithms is with Rolf Radespiel of DLR, the German analog of NASA. Rolf's group is developing a Navier Stokes code for hypersonic flows in collaboration with NASA Langley. One problem occurring in such codes is slow convergence of the multigrid algorithm due to the highly stretched grid needed to resolve the boundary layer. We anticipate that techniques analogous to those used in the elliptic case can resolve this problem, and are thus in the process of modifying the DLR/NASA hypersonic code to implement one of our algorithms.

Having robust algorithms requiring only point relaxation is important for Navier Stokes code, since line- or plane-relaxations can be prohibitively expensive for systems of equations. The particular algorithm we are using is based on coarsening the grids by a factor of four in each of the coordinate directions, rather than the usual factor of two. The result of this faster coarsening is that the total cost of the algorithm is about twice the cost of the fine grid iteration, despite the large number of “semicoarsened” grids which occur.

Computational Aerodynamics

Dimitri Mavriplis

Work is continuing on the use of unstructured meshes for solving computational fluid dynamics problems in both two and three dimensions.

In two dimensions, a new effort has been initiated in the area of grid generation. A grid generation methodology is being developed in order to enable the automatic generation of meshes about arbitrary two-dimensional geometries for viscous flows, with the possibility of producing highly-stretched elements in boundary layer and wake regions. Concepts which are successfully proven in two-dimensions will eventually be extended to the three-dimensional setting.

In three dimensions, the previously developed unstructured multigrid Euler solver has been employed to compute the flow over a fuselage-wing-nacelle aircraft geometry using a fine mesh of 804,000 vertices, and 4.5 million cells. This is believed to constitute the largest unstructured grid problem attempted to date. The choice of such a problem was designed to demonstrate the efficiency of the unstructured multigrid approach. The same case was run on the CRAY-YMP machine, using all 8 processors, the CRAY-YMP-C90 machine using 16 processors, and the Intel Delta iPSC860 massively parallel machine, using 512 processors. A computational rate of 750 Mflops was achieved on the CRAY-YMP, while a rate of 3.2 Gflops was achieved on the CRAY-YMP-C90. The rate observed on 512 Intel Delta processors was of the order of 1.5 Gflops. The steady-state solution for this case was achieved in 100 multigrid cycles, within which the residuals were reduced by six orders of magnitude. This required a total of 268 seconds on the fastest machine, the CRAY-YMP-C90.

The results achieved on the Intel Delta machine were the result of an on-going effort aimed at studying the numerical and computer science issues of mapping irregular problems to massively parallel distributed memory architectures. This work is being carried out in conjunction with J. Saltz, Subhendu Das, Ravi Ponnusamy, and Satya Gupta, using a set of software primitives known as PARTI. This work is aimed at demonstrating the applicability of the PARTI primitives to practical implementations, and also to employ experience gained

from particular implementations to aid in the design or modification of the primitives, as well as their potential influence on the design of a prototype parallelizing compiler. The most recent results of this effort have been presented at the 30th Aerospace Sciences Meeting and Exhibit in Reno Nevada, in January 1992, as AIAA paper 92-0562 (also ICASE report 92-12). The results achieved on the CRAY-YMP and CRAY-YMP-C90 machines will be presented at the Supercomputing '92 meeting, to be held in Minneapolis, Minnesota, in November 1992. In the future, research in parallel computations will be extended to include operations such as unstructured mesh generation, parallel partitioning, and the use of adaptive meshing strategies.

James Quirk

Work is continuing on a cartesian grid scheme for simulating shock hydrodynamic phenomena that involve awkward geometries. Thus far the scheme has been shown capable of handling arbitrarily complex, two-dimensional bodies (ICASE Report 92-07), provided that the bodies remain stationary. Currently my efforts are concentrated on trying to extend the scheme so as to allow for moving bodies.

Paul Saylor

In work done in association with Jeffrey Scroggs, a method of lines approach is used to solve the transonic small disturbance equation. For this purpose, the package DASSL, due to L. Petzold, is being used. DASSL is a general solver for differential algebraic equations, which uses implicit schemes and has been modified for the the solution of large scale non-linear systems. Besides the implicit formulation, other advantages to this approach are: the automatic step size control; and high order temporal accuracy. Formulation of the airfoil and wake conditions has been done with the help of LaRC staff.

There is a large number of nonsymmetric eigenvalue problems that arise in numerical studies of transition to turbulence. Recent algorithmic development makes it possible to consider kernel polynomials to enhance midrange eigenvalues. This approach is related to ideas on polynomial filtering described in Sorenson (SIAM Journal on Statistical and Scientific Computing, 1992) but is more closely related to filtering ideas traceable to Stiefel. The flexible nature of kernel polynomials is desirable in the midrange.

Jeffrey Scroggs

This work is in close collaboration with Paul Saylor, and also in collaboration with Steven F. Ashby (Lawrence Livermore National Lab). In it, we combine asymptotic analysis and

domain decomposition to derive preconditioners for use in the iterative solution of linear systems associated with singularly-perturbed partial differential equations. The analysis is used to obtain the first term in a uniform asymptotic expansion of the solution. This requires a “natural” decomposition of the computational domain in which each subdomain exhibits a fundamentally different type of physical behavior. One or more terms of the original PDE are then identified with this behavior, and these “partial” operators are individually discretized and assembled into a preconditioner. Several variants of these “physically motivated domain decomposition” preconditioners are compared to standard preconditioners in the GMRES method on a linear two-dimensional singularly-perturbed convection-diffusion equation.

Shlomo Ta’asan

The development of efficient multigrid solvers for constrained optimization problems governed by partial differential equations has continued with research in two directions. The first, which is well developed by now, deals with problems in which the parameter space on which optimization is done is of finite dimension in the differential formulation of the problem. The methods use relaxation for the parameter space in a multilevel way. Parameters that have a non-smooth effect on the solution are relaxed on fine levels while those of smooth effect are solved for on coarse grids only. The methods use adjoint variables to define a descent direction for the minimization problem. The other direction focuses on problems in which the optimization is over an infinite dimensional parameter space (in the differential level). Also here the same type of ideas for the treatment of the different scales in the problems is being used. Experiments with some model problems involving elliptic partial differential equations as the constraint equations have been performed showing that the full optimization problem can be solved with a computational cost which is only a few times more than that of solving the PDE alone.

The above ideas have been applied in aerodynamics design problems where airfoils are to be calculated so as to meet certain design requirements, for example, to give pressure distribution in some flow conditions which are closest to a given pressure distribution. The present model for the flow is the transonic full potential equation with a body fitted grid. The shape of the airfoil in these calculations is being expressed in terms of a finite number of given shape functions with amplitudes to be found by the design process. Currently, subsonic design problems are being investigated. The goal is to obtain a solution of the optimization problem in a computational cost which is just a few times (2-3) that of the flow solver. Such a result was demonstrated for a simpler model of the flow. This work is jointly done with M. D. Salas (Fluid Dynamics Division, LaRC) and G. Kuruvila (Vigyan, Inc.).

New multigrid solvers for inviscid flow problems are being developed in which the conver-

gence rates are to be like that of the full potential equation. These convergence rates have been predicted by Fourier analysis. These solvers employ relaxation methods of the Gauss-Seidel type with a proper modification to handle systems of partial differential equations. At present we focus on the incompressible and compressible inviscid case working with body fitted grids in two dimension, using MAC discretization schemes. Flow over a cylinder is being investigated.

Bram van Leer

Three topics were addressed:

- Application of local preconditioning to the Euler equations in order to remove stiffness due to large a cell-aspect ratio. It appears that this is not entirely possible, and that non-local preconditioning techniques are needed in order to handle the stiffness caused by acoustic waves moving streamwise and counter-streamwise.
- Development of a boundary conditioning simulating a perforated wall insert. This is needed for Navier-Stokes simulations done in support of shock-boundary-interaction experiments by Rosemary Rallo (Exp. Flow Phys. Branch).
- Development of approximate Riemann solvers for systems of equations that are not entirely in conservation form. These arise, for instance, in generalized magneto-hydrodynamics, where the electrical potential energy of ions and electrons is preserved in total, but not for each species separately.

Control Theory

Kurt Bryan

Work continues with W.P. Winfree of the Nondestructive Evaluation Sciences Branch, LaRC, on a thermo-mechanical inverse problem using the response of a structure to external heating to determine its internal properties and locate defects. The recent focus has been on using the boundary temperature response of the sample to recover information about internal defects (voids). The heat conduction problem is formulated as a system of boundary integral equations. Uniqueness and continuous dependence results have been proved for this thermal inverse problem, and the boundary integral equation formulation has been implemented numerically in an optimization/ fit-to-data approach for solving the inverse problem. Convergence results for this approach to the inverse problem have been proved. Numerical simulations have been encouraging. We are now pursuing a sensitivity analysis as well as determining optimal design parameters for experimental work and should begin collecting actual data soon.

Work is also continuing with Michael Vogelius of Rutgers University on the inverse conductivity problem, another imaging technique which finds use in the nondestructive evaluation of materials. The goal is to recover a collection of several internal cracks in an electrical conductor based on voltage and current flux measurements on the surface. A numerical algorithm for the recovery of a collection of linear cracks has been developed. This procedure also makes use of an integral equation formulation for the forward problem. The iterative algorithm, based on Newton's method, adaptively changes the applied current flux at each step in order to achieve maximum sensitivity to the location of the cracks. The numerical simulations here have also been extremely encouraging, thus validating this approach to the inverse problem.

Fumio Kojima

In collaboration with H.T. Banks and W.P. Winfree (Instrument Research Division, LaRC), work is continuing on geometrical heat inverse problems arising in thermal tomography. We are developing the software for identifying structural flaws inside a body using thermographical data. Three different approaches for attacking that identification problem have been investigated. One approach is a spline-based method for estimating corrosion profiles of the back surface using the data from the front surface. Major efforts on this approach during the past several months were focussed on the robustness analysis of identification algorithm with respect to data perturbations. A graphical scheme was considered for adjusting the "optimal" regularization parameter related to the Tikhonov's method. Although there are still many open problems on the theoretical part of this approach, the validity of the scheme proposed were demonstrated through computational experiments and experiments with laboratory data. A second approach was tested for the detection of bonding failures in aircraft lapjoints and in the adhesive joints between aircraft skin and reinforcing doublers. A fast algorithm was proposed using the domain decomposition method. We are currently examining the proposed algorithm with laboratory data. A third approach for assessing the structural integrity of materials is a learning algorithm using multi-layer feed-forward neural networks. The trust region algorithm is adapted to the backpropagation supervised learning problem. Training and testing were carried out effectively using the laboratory data. The development of the parallel algorithm for high dimensional problems is underway.

Yuh-Roung Ou

In a joint effort with John Burns, work is continuing on the development and analysis of mathematical models for use in control of fluid dynamic systems. This work is aimed at developing theoretical tools and numerical algorithms to solve the necessary conditions

for the complete optimal flow control problems. The method newly developed here will be used to investigate fundamental questions regarding control of separated flows by using various boundary control mechanisms (i.e. moving surfaces, blowing and suction). The primary focus of our attention so far has been on flow simulation around a rotating cylinder. A computational scheme has been developed for the purpose of determining the control (i.e. rotation rate) that maximizes the lift-to-drag ratio. This moving surface mechanism is demonstrated to be intimately relevant to the vortex shedding process in the wake and development of forces acting on the cylinder surface. In the computation, all the optimal solutions are directly obtained by trial and error variation of controls. However, our future effort will focus on the feedback control problem that has not been resolved.

A joint research involving mathematical theory, numerical computations and laboratory simulation on optimal control problems associated with viscous flow past a cylinder was conducted in collaboration with S. S. Sritharan (University of Colorado) and E. Hendricks (Naval Ocean Systems Center). The main mathematical results are an existence theorem for the optimal control and dynamic programming using the infinite dimensional Hamilton-Jacobi-Bellman equation to provide the feedback map. An experimental study into the use of a pair of suction/blowing slots at cylinder boundary in preventing unsteady boundary layer separation is underway. Numerical computations are performed for the optimal control of time-averaged forces on a rotating cylinder utilizing rotation rate as a control parameter. Further numerical investigation will focus on implementation of a computational algorithm to calculate the optimal solution based on the theoretical approach.

Ralph Smith

In collaboration with H.T. Banks, work is continuing on the development of a noise suppression methodology for structural acoustics problems. This involves the modeling of underlying physical phenomena, the identification of unknown physical parameters through data fitting techniques, and the development of feedback control techniques for reducing noise which is being generated by the vibrations of an elastic structure.

The general model under consideration consists of an exterior noise source which is separated from an interior chamber by an elastic boundary (for example, a thin beam, plate or shell). This bounding structure transmits noise or vibrations from the exterior field to the interior cavity via fluid/structure interactions thus leading to the formulation of a system of partial differential equations consisting of a 3-D acoustic wave equation coupled with structural elasticity equations. Control is implemented via piezoceramic patches which are bonded to the boundary of the acoustic cavity. In this way, we can take advantage of the natural "feedback" loop which is due to the coupling of the structural vibrations and their

radiating or receiving acoustic fields. The validity of the model is being tested both numerically and with data from experiments by H.C. Lester and R.J. Silcox (Acoustics Division, LaRC).

As a first step toward developing an effective linear quadratic regulator (LQR) time domain state space control methodology for near field acoustics problems of this type, we have considered a simplified but typical linear model consisting of a 2-D interior cavity with an active beam at one end. Numerical tests have shown that when the feedback controls are implemented, substantial reductions in acoustic pressure can be realized with physically realistic voltages into the patches. Current work on the 2-D problem is centering around the identification of physical parameters and the development of LQR control techniques for the more general model involving nonlinear coupling terms.

Current and future work is also being directed toward the control of noise in 3-D structural acoustics models. In the first such model, the interior cavity is taken to be cylindrical with a circular vibrating end plate and sectorial patches. An accurate, efficient and stable means of discretizing the resulting system of coupled PDE's has been developed and efforts are being directed toward the control of acoustic pressure through techniques similar to those developed for the 2-D problem.

In the second 3-D model of interest, we assume that the cylindrical surface itself is vibrating and thus causing unwanted interior cavity noise. In order to develop an effective feedback control scheme using piezoceramic actuators, the interaction between the patch and the underlying cylindrical shell must first be accurately modeled. The patch in this model is assumed to be curved and the coupling between the bending and in-plane deformations, which is due to the curvature, is retained. Once this modeling phase has been completed, the dynamic piezoelectric patch/cylindrical shell model can then be coupled with an acoustic pressure response inside the cylinder in a manner similar to that used in the earlier models. The problem in this format is then amenable both to the identification of physical parameters through data fitting techniques and to optimal control techniques analogous to those developed for the 2-D model.

FLUID MECHANICS

Fabio Bertolotti

Two projects are currently being investigated. The objective of the first project is the development of a family of codes for the analysis of transition over swept wings of infinite span in subsonic flows. The core program will be a PSE solver written in surface-oriented coordinates, and employing a multidomain spectral method in the surface normal direction coupled to a finite difference method in the streamwise direction. The use of primitive variables will allow for a rapid extension to compressible flow. An accurate mean-flow solver, similar to the PSE solver, has been developed. The calculations, which employ a new set of governing equations, solve simultaneously for the viscous (i.e. boundary layer) and potential flow. Consequently, no Goldstein singularity exists at separation, and the solver can streamwise march through a small separation bubble.

The objective of the second project is the analysis of nonlinear mode interaction in a subsonic wake. This work is being performed in tandem with experimental investigations by Tom Corke at the Illinois Institute of Technology. In contrast to boundary-layer flow, the wake is absent of wall regions with strong shear, thus altering the strengths of the modal interactions. The close cooperation with experimental investigations may lead to a more complete road-map to the transition process in wakes.

Peter Duck

The study of the effect of disturbances on the shock wave on a wedge in a supersonic flow is continuing. These disturbances can be in the form of vorticity waves, entropy waves or acoustic waves. These waves are then modified upon passage through the shock, creating a complex disturbance flow field. The effect of the disturbance on the boundary layer on the wedge is also being studied.

The stability of compressible vortex flows is also being investigated. Here a vortex flow is perturbed, and the (inviscid) stability is determined. This is being considered both numerically, and also using asymptotic techniques valid in the limit of large Mach numbers.

Gordon Erlebacher

Work on shear flow compressible turbulence in collaboration with Sarkar, and Hussaini has been completed. The analysis of the compressible shear flow simulations of Sarkar and myself has been completed. A paper has been written and submitted to Physics of Fluids.

This project brought out some of the effects of compressibility on the small scale structures in compressible homogeneous shear layers.

Work is in progress with Craig Streett. The objective of this project is to develop a consistent formulation for the stability problem on cones in the hypersonic regime, both at zero and at finite angles of attack. That analysis hinges on the choice of a preferred coordinate system. At this stage, a spectrally accurate mean flow code, including the shock, has been developed and tested. Existing linear and secondary stability codes have been adapted to this problem and are currently undergoing testing.

Work is in progress with Shu (Brown University) on the interaction of a shock with an impinging longitudinal vortex. This work has relevance for the stability of aircraft. Both computational and theoretical aspects of this problem are under investigation. Both axisymmetric and 3-D codes have been implemented using a 3rd order ENO scheme.

Work continues with Biringen at Colorado State University on the Direct Numerical Simulation of compressible Couette flow.

James Geer

Several variations of a hybrid perturbation/variational technique, which have been under development for some time, are being applied to a variety of different problem areas. In particular, the technique is being used to compute semi-analytical, semi-numerical approximations to the resonant frequencies for Hamiltonian or Lagrangian systems of nonlinear ordinary or partial differential equations. In the context of this problem, a variety of different formulations of the technique are being compared to determine the relative merits of each approach. The technique is also being applied to several classes of elliptic boundary value problems with irregularly shaped domains. The basic idea here is first to use some simple homotopy ideas to embed the problem of interest in a family of problems, which can be analyzed using standard perturbation techniques. Then, the perturbation solutions are “improved” using a variety of Galerkin type approximations.

A three-step hybrid analysis technique, which successively uses the regular perturbation expansion method, the Pade’ expansion method, and then a type of Galerkin approximation, is also being developed and studied. It is being applied to several model problems which develop boundary layers as a certain parameter becomes large. In particular, the technique appears to simulate these boundary layers by producing approximate solutions with singularities which lie just outside the domain of interest. Based on some preliminary results, the technique appears to provide a good approximation to the solution, even when the perturbation and Pade’ approximations fail to do so.

Another new technique (temporarily called the “flexible singularity expansion technique”)

is being developed to determine a family of approximate solutions, based on a small parameter ε , to certain classes of exterior boundary value problems. These solutions are expressed in terms of singular solutions to the governing differential equation for which the singularities lie outside the region of interest. In general, the exact type, location, and strength of these "flexible" singularities are determined by the governing equation and shape of the domain of the problem. More specifically, the various parameters associated with these new singularities are determined by requiring that the "flexible singularity solution" agree with the perturbation solution to the problem to within a prescribed order in ε as $\varepsilon \rightarrow 0$. The technique is being applied to several classes of problems involving nonlinear PDE's and to problems which are geometrically nonlinear.

Chet Grosch

Research on two separate topics is being pursued. The first topic is the interaction of a disturbance wave with a shock, taking into account the internal structure of the shock. The interaction of a sound or vorticity wave with a shock has been studied in the past using the approximation that the shock is an infinitesimally thin discontinuity. In the current study the shock is taken to have a non-zero thickness. The internal structure of the shock is determined by solving the compressible Navier-Stokes equations assuming that the gas is perfect and that the viscosity-temperature relation is given by the Sutherland relation. The wave disturbance which interacts with the shock can be either a sound wave or a vorticity wave. These disturbances are generated upstream of the shock and propagate into the shock. The object of the study is to determine the "transmission coefficient" of the shock as a function of the shock and wave parameters. It is also planned to extend this study to the case of a shock initiated detonation wave. The research is being carried out in collaboration with D.G. Lasseigne.

The second topic being studied is the stability of reacting, compressible shear flows which are not strictly parallel. Thus the standard parallel flow stability may not be an adequate approximation. A case in point is the stability of such a flow in the ignition region where the flow may be changing rapidly with downstream distance. In order to account for these, and similar effects, it is necessary to develop a theory for the non-parallel correction to the stability calculations for reacting compressible flows. The theory has been developed and calculations for specific reacting flows is under way. The research is being carried out in cooperation with D.G. Lasseigne and T.L. Jackson.

Philip Hall

The nonlinear evolution of the fastest growing Görtler vortex mode was considered using a combined finite difference-spectral approach. It was shown that the vortices cause the sublayer in which they grow to break down and flow separation to occur. Results on the effect of wall-cooling and gas dissociation on Rayleigh instabilities were obtained. In addition the effect of Görtler vortices on Rayleigh modes was discussed. An investigation of the absolute/convective nature of small wavelength vortices was carried out. Work on heated boundary layers has shown the importance of streamwise vortices in initiating inviscid secondary instabilities.

Tom Jackson

Work focuses on different aspects of flame/vortex interactions, a fundamental problem for the understanding of small scale structures in turbulent combustion. Particular aspects under consideration include acoustic influences, compressibility effects, and thermal expansion. A combination of asymptotics and numerics will be used to reduce the complex problems to a model problems, thus isolating key physical effects for analysis. This work is in collaboration with M. Macaraeg (Fluid Mechanics Division, LaRC) and M.Y. Hussaini.

D. Glen Lasseigne

Research on a number of topics is being pursued. First, the oblique detonation attached to a wedge is being investigated for stability and its response to imposed disturbances. Secondly, we are investigating the interaction of a sound wave with a shock taking into account the internal structure of the shock, which is taken to have a non-zero thickness. Finally, the theory of the non-parallel correction to the stability of a reacting supersonic shear layer is being developed.

Stephen Otto

Current interest is concerned with the solution of the equations governing the interactions between a fully nonlinear Görtler vortex structure and a linear inviscid Rayleigh wave. The flow situation being considered is that of an imposed pressure gradient on the flow over a plate of variable curvature. This is joint work with P. Hall, and is considered in both the incompressible and compressible flow regimes. Also underway is a speculative study of vortex wave interactions for vortices of order one wavenumber. This work is aimed at explaining the sudden growth of a vortex structure in the numerical experiments of Streett & Joslin. Work is underway with J. P. Denier (Univ. New South Wales, Australia) to consider the

fully nonlinear development of vortices occurring in a Rayleigh layer occurring at a cylinder undergoing spin up.

It was found in Denier, Hall & Seddougui (1991) that there may exist modes with wavenumbers of $O(G^{\frac{1}{2}})$ where G , the Görtler number, is taken to be large. These modes have been dubbed the most dangerous Görtler modes, due to their possessing a growth rate of the order of $G^{\frac{3}{2}}$. It was found that by the imposition of crossflow that it is possible to stabilise these modes. The linear analysis of Bassom & Hall (1991) has been extended to the weakly nonlinear regime, and work is currently underway with A. P. Bassom (Univ. of Exeter, UK) to extend this analysis to a fully nonlinear regime.

Sutanu Sarkar

We are primarily engaged in the direct simulation and Reynolds stress modeling of turbulent flows. Continuing our investigation of compressibility effects in shear flows, we find that the pressure field is significantly altered from the baseline incompressible case, thus altering the Reynolds stress anisotropies. The collaborative study with G. Erlebacher of the eigenstructure of the rate of strain tensor and orientation properties of the pressure gradient vector and vorticity vector, has revealed both similarities and a few important differences with the incompressible case. The performance of advanced Reynolds stress closures in the log-layer and wake region of incompressible channel flow has been studied in collaboration with D. Demuren in order to assess their strengths (relative to simpler two-equation models), and to identify and correct deficiencies. We have begun an approach in computational aero-acoustics which couples the calculation of the unsteady flow field in a homogeneous turbulent flow via DNS of the compressible Navier-Stokes equations with an appropriate acoustic analogy in order to compute the far-field sound.

Charles Speziale

An examination of methods for testing models for the pressure-strain correlation of turbulence using DNS data bases has been conducted in collaboration with S. Sarkar and T.B. Gatski (Fluid Mechanics Division, LaRC). It was shown that in homogeneous turbulent flows with constant mean velocity gradients, the decomposition of pressure-strain models into slow and rapid parts is ambiguous. Consequently, when pressure-strain models are decomposed into slow and rapid parts along the traditionally assumed lines – and separate comparisons are made with DNS data bases for these correlations – highly misleading conclusions can be drawn about the predictive capabilities of a model. Other pitfalls in the testing of pressure-strain models were also uncovered and more reliable tests were proposed.

Work has continued with T.B. Gatski on the development of practical models for the anisotropy of dissipation based on an extension of a modeled transport equation for the dissipation rate tensor that was developed earlier (ICASE Report No. 90-88). By means of a local-homogeneous-equilibrium hypothesis, algebraic models for the anisotropy of dissipation can be obtained in a systematic fashion. These algebraic models – which unlike previous models depend nonlinearly on the mean velocity gradients – are solved in conjunction with a derived transport equation for the scalar dissipation rate wherein the coefficient of the production of dissipation depends on the invariants of the mean rotational and irrotational strain rates. The new model is currently being tested in a variety of benchmark homogeneous turbulent flows. Research along these lines appears to have the potential to lead to a new generation of two-equation models and second-order closures.

Research has also been conducted with T.B. Gatski and R.M.C. So (Arizona State University) on the development of a near-wall compressible second-order closure model. Major numerical stiffness problems must be overcome before second-order closure models can be integrated directly to a wall in turbulent supersonic flows. This is the ultimate goal of the present study.

Experimental Research

Thomas Corke

This work is aimed at mechanisms which govern the growth of secondary three dimensional modes of a particular type which feed from a resonant energy exchange with the primary (most amplified) two dimensional mode in a Blasius boundary layer, and in the far wake of a 2-D body. Our approach was to introduce controlled time periodic 3-D (oblique) wave pairs of equal but opposite sign, simultaneous with a 2-D wave. The waves were introduced by an array of v-component producing elements. In the case of the boundary layer, the relatively lower frequencies of unstable TS modes allowed the use of heating wires, elevated from the wall at the height of the critical layer. In the wake, metallized electrodes which were vapor deposited onto a piezoelectrically active polymer wrapped around the surface provided the disturbance input. The amplitudes, streamwise and spanwise wave numbers, and initial phase difference are all individually controllable. Comparisons between the two flowfields show similarities and interesting contrasts. The initial work in the wake has focused on 2-D fundamental/3-D subharmonic sinuous mode interactions, and the dependence on spanwise wave number. In the boundary layer, detuned fundamental/subharmonic mode resonance has also been investigated. In the relatively weak mean shear present in the far wake, the resonant growth of the subharmonic 3-D mode was followed downstream by an exclusive pe-

periodic growth and decay of the primary 2-D and 3-D modes. In the boundary layer, the 3-D mode gains energy from the strong mean shear, leading to its dominant growth downstream. Energy transfer to other sum and difference modes leads to gradual spectral broadening and transition to turbulence. The results include second and third order cross-spectral analysis, mode eigenfunction modulus and phase distributions in space; and stream function, vorticity and particle tracer distributions for the phase reconstructed flow fields.

COMPUTER SCIENCE

Tom Crockett

The primary activity was preparing the iPSC/860 complex for transfer from ICASE to Langley's Analysis and Computation Division. A detailed transition plan was developed which maintains functionality for the existing user community while eliminating dependencies on ICASE machines for system services. Implementation of the transition plan began, and ACD personnel are being trained to take over system administration duties. To facilitate the transition, center-wide conventions for assigning user IDs and naming filesystems were developed in conjunction with representatives from ACD and the Langley distributed computing community.

Beta test versions of several Intel software components were installed, and many trouble reports and feature requests were delivered to Intel.

Planning work for a scientific visualization research activity at ICASE began. Infrastructure requirements were identified, and a near-term spending plan was developed.

Design and implementation of a parallel graphics library began. This library will incorporate the Orloff/Crockett rendering algorithm and will include additional functionality and performance improvements needed by visualization applications running on parallel systems such as the iPSC/860 and CM-5.

Raja Das

We are developing techniques to efficiently parallelize irregular problems. Based on these techniques we have built a set of tools that can be used directly by users or by compilers to parallelize irregular loops. Using the tools we have parallelized a 3-D unstructured multi-grid Euler equation solver and also parts of the molecular dynamics code CHARMM. The Euler solver is running on both the iPSC/860 and the Touchstone Delta machine. Parallel CHARMM has been run on the iPSC/860.

We have also developed a prototype compiler to perform irregular loop transformations. The compiler takes as input fortran programs with Fortran-D like data distribution statements and generates code that can be executed on a parallel machine. In the parallel code, all irregular loops are split into two constructs; the inspector and the executor. The inspector part of a loop does the required preprocessing and the executor does the actual computation. Both the inspector and the executor is built using the tools we have developed. This research is done in collaboration with Joel Saltz and Dimitri Mavriplis.

We plan to develop tools and techniques to link partitioners to parallelizing compilers. This work will be based on the infrastructure that exists for the Fortran-D project at Rice Uni-

versity. We are also planning to develop techniques that can be used to parallelize adaptive irregular problems. We will use these tools and parallelize a simple mesh refinement program and also finish the parallelization of CHARMM. This work will be done in collaboration with Joel Saltz.

Satyanarayan Gupta

A version of Parti has also been developed for the Maspar MP-1. The primitives developed for the MASPAR using the same user interface as that of MIMD primitives for ease of portability and development of compilers for architecture independent parallel languages such as Fortran-D and Fortran-90. The implementation and optimizations carried out in these primitives are different than those adopted for the MIMD machines. The optimizations carried out to improve communication performance on the Maspar include 1) removal of duplicate off-processor accesses, 2) partitioning to increase ratio of intra to inter-cluster communication and to reduce communication volume, and 3) preprocessing to increase router utilization. During each router cycle, only one processor from each cluster can send data and only one processor can receive data. Thus for each cycle of the communications schedule, we pair up processor clusters and decide in advance which processors will transmit data to which other processors. Mavriplis' 3-D unstructured Euler solver has been ported to the Maspar using earlier versions of Maspar Parti. It appears that the use of all of these optimizations together will result in a performance improvement of roughly a factor of 7 relative to the initial naive implementation. This is joint work with Joel Saltz

David Keyes

There are three paradigms for partitioning the work of solving a partial differential equation system such as Navier-Stokes for parallel computation: splitting up the operator(s), splitting up the solution space, and splitting up the domain. None of these should be overlooked, but it is the third — domain decomposition — that seems to best simultaneously fit both a law of nature, that influence decays with physical distance, and a law of distributed computation, that data costs increase with network distance. Exceptions exist, and other paradigms may possess elegance and convenience so that they should be used in combination with domain decomposition, either inside or outside. In general, however, conventional operator and function space decomposition methods are more limited in scalable parallel performance because they tend to reach for less important and more expensive data more frequently.

For this reason, we have been exploring the capabilities of iterative domain decomposition algorithms on what might be termed model engineering problems: Navier-Stokes at low

Reynolds number in simple geometry, and combustion in simple kinetics limits. We have developed extensions of the mathematical theory for convergence of domain decomposition algorithms to nonsymmetric and/or indefinite systems and have collaborated in the implementation of domain decomposition codes on a variety of parallel computer platforms for performance testing. Recent work has identified Krylov-accelerated multiplicative Schwarz methods as very effective in the highly convective regime. Tests on higher Reynolds number and transonic Mach number situations are in progress.

Piyush Mehrotra

Large scientific applications, such as multi-disciplinary optimizations codes, exhibit several types of parallelism. We have been studying an object oriented methodology to exploit both functional parallelism at the outer level and data parallelism at the inner level. In our approach, objects are used both for expressing parallelism and for data encapsulation. Control and functional parallelism, exhibited at the top level by such applications, can be captured by asynchronously executing objects. On the other hand, data parallelism within each object can be expressed through one of the data parallel languages. In contrast to most other parallel object oriented systems which only support only one level of parallelism, our approach exploits parallelism both across and within objects. Objects can, therefore, be thought of as process groups interacting with each other through distributed method calls.

One of the underlying goals of our system is to provide support for heterogeneous systems in a transparent manner. That is, the interaction of two objects should be expressible independent of where the objects may be executing. The objects could both be on the same parallel machine or on different systems communicating across a local area network. As far as possible, we would like to exploit existing technology and systems for our implementation. Thus, we are considering systems which support distributed computing such as PVM from Oakridge and Hermes from IBM as a core around which to build our prototype system.

In a joint effort with Prof. Hans Zima and Barbara Chapman of University of Vienna, we have been developing a language, called Vienna Fortran, for expressing data parallel algorithms. Our approach, though similar to that followed by other researchers in the field, differs from these efforts in significant ways. The main aim of all these approaches is to allow the user to control the distribution of data while still using a global name/index space for the code. Vienna Fortran provides an extensive set of mechanisms to distribute and align arrays including intrinsic and user-defined functions. A clear distinction is made between arrays which are statically distributed and those whose distribution can be changed at runtime. This not only allows the compiler to make optimizations for the statically distributed arrays. but also helps the user in distinguishing between the two kinds of arrays.

Vienna Fortran provides a complete set of extensions for Fortran, paying attention to such Fortran issues as sequence and storage association, the distribution of arrays in common blocks, and distributed arguments for procedures (see ICASE Internal Report 21). A forall loop (with reduction operators) allows the user to express parallel iterations. The concept of allocatable arrays (borrowed from Fortran 90) allows the use of distributed work arrays. The language also includes facilities for concurrent I/O of distributed data structures. A compiler for the language is currently being implemented at the University of Vienna.

In work done in close collaboration with John Van Rosendale, it was found that in recent years a number of languages have emerged for expressing “data parallel” algorithms for distributed memory architectures. The authors’ language, Kali, was one of the first of these, borrowing “forall loops” and reduction operators from Sisal and Blaze, and also adding a number of new constructs, such as array distribution clauses, specific to distributed memory machines. Together these constructs allow one to express numerical algorithms for distributed memory architectures about as easily as one expresses sequential algorithms.

Kali has since been joined by a number of siblings, including Dyno, Fortran D, and Vienna Fortran. All are semantically similar, sharing with Kali the fundamental constructs described, but differing in many small and perhaps not so small ways. It is thus natural to begin looking at the expressive power of these languages, and the extent to which they suffice for expressing complex scientific calculations.

These Kali-like languages are designed for regular computation, the model being dense linear algebra or Jacobi-style iteration on rectangular grids. Most of these languages can also readily express more general “uniform” computations, such as tensor product algorithms. Thus the limitations of these languages show up in algorithms requiring either irregular data structures, or asynchronous (MIMD) computation, or both. Our present focus of study is on algorithms based on irregular data structures.

One limitation of this class of languages has already manifested itself in the context of iterative methods on irregular meshes, such as the Euler and Navier Stokes solvers being designed by Dimitri Mavriplis. In many of these programs one has several related data structures. For example, one might have lists of both edges and vertices. In this case, both edges and vertices need to be distributed across processors, and furthermore, these distributions should be aligned as well as possible, to avoid excess interprocessor communication.

A different kind of problem arises in “partially structured” computations, such as that occurring with block structured grids. Kali-like languages support both structured and unstructured computations, but are naturally more efficient for structured computations. With a block structured grid, for example, one would like to fully exploit the structure in each block, yet support completely general inter-block computation. Precisely analogous

issues arise for multigrid algorithms.

The result of this effort will be, at minimum, a careful critique of Kali-like languages for distributed memory architectures. There is a real need to know their range of applicability, and when can they be used effectively. In general, we feel that the present limitations of these language are not intrinsic, and can be readily corrected. For example, one can easily add a “message passing” layer on top of a Kali-like language. In some cases, this could dramatically increase the expressive power of these languages.

David Nicol

A number of numerical algorithms require frequent, global redistribution of data. In such situations it is important to be able to execute a global load balancing as quickly as possible. Previous approaches to the problem suffer either from non-scalability, or from moving more data than is necessary to accomplish the balance. We have developed new scalable algorithms which are optimal in the sense of not moving more data than necessary to achieve balance. Studies of the algorithm on simulated large scale systems prove its superiority over existing methods.

In work done in collaboration with Albert Greenberg and Boris Lubachevsky (AT & T Bell Labs), we are developing parallel MIMD and SIMD algorithms for the discrete-event simulation of networks where messages may be re-routed in the presence of network congestion. We have implemented the algorithm on the Intel iPSC family of multiprocessors, and have achieved a goal of simulating over one million messages routed per second, on 64 nodes of the Touchstone Delta machine.

Also, in work done in collaboration with Phil Heidelberger (IBM Research), we are developing new synchronization algorithms for parallel simulation, to be used when the underlying system is a Markov chain. Our approach is based on the method of uniformization, which essentially gives us a mechanism by which we can predict exactly when communications occur. We have implemented both conservative and optimistic versions of the method, achieving speedups of at least 180 on 256 Touchstone Delta nodes.

Terrence Pratt

A promising new method, called the KCD method, for automatic generation of an MIMD, distributed memory, parallel program from existing sequential code is being studied. The method is based on a theoretical model called “kernel-control decomposition” (KCD). The project has developed a new method for decomposition of the source program into pieces appropriate for parallel execution and a new model for parallel program execution based on

this decomposition. In addition, a new technique for generating close-to-optimal overlap of communication and computation during program execution has been developed.

These concepts have been implemented in a software environment named PISCES 5. Using this software, we are exploring the application of these methods to applications codes, such as the LAURA code developed by Peter Gnoffo (Space Systems Division, LaRC) and the thirteen codes in the Perfect Club benchmark suite. Measurements of the performance of the parallel versions of these codes are being made on the 32-node NASA/ICASE Intel iPSC/860.

Matthew Rosing

Work in developing a low latency message interface and language and compiler techniques for the efficient use of distributed memory multiprocessors continues.

We are developing a set of communication primitives that are implemented much more efficiently than the send/receive message model currently supported on distributed memory machines. The send/receive model is the primary reason why the software latency of transmitting a message on newer machines is around 50 microseconds, even though the hardware component is around 100 nanoseconds. This imbalance makes it difficult to write a large class of programs that naturally have short message lengths such as adaptive and irregular algorithms. In support of this goal, on the ipsc/860, we have been able to reduce the software latency, in many cases, to only a few machine instructions. This work is being done in conjunction with J. Saltz.

I am working on language and compiler techniques for writing libraries for numerical applications that are highly efficient, flexible, and easy to use. Due to the inability of libraries to adapt to the way in which they are used, they tend to be very inflexible and therefore inefficient. In response to this, we are developing techniques to build libraries containing procedures and functions that can adapt to the context in which they are used. Although these procedures are not libraries in the traditional sense, they appear that way to the end user. This work is based on our experience with writing compilers and is being done in conjunction with R. Schnabel at the University of Colorado.

Joel Saltz

We have developed a suite of primitives that track runtime data dependencies. These primitives employ a static task and communication schedule on each processor. A preprocessing phase determines the precise sequence in which work and communication are to be carried out. These primitives have been used to port sparse direct methods to the iPSC/860.

Results are very favorable relative to hand coded sparse factorization methods. The absolute performance results for these and other distributed memory implementations of sparse factorization methods are still unimpressive and we are actively seeking improved performance. Longer term, these primitives will evolve into the runtime support needed to handle problems with runtime determined loop carried dependencies such as sparse factorizations and sparse triangular solves. This work is joint with Sesh Venugopal (Rutgers University) and Vijay Naik (IBM Hawthorne).

The Parti primitives can be used directly by programmers or incorporated into compiler for distributed memory architectures. They are currently being used to implement and extend Fortran D. Joint work with Raja Das, and with Ken Kennedy, Chuck Koelbel and Rinehard von Hanxleden at Rice; Geoffrey Fox and Alok Choudhary at Syracuse (NSF Center for Research on Parallel Computing).

Fortran and C callable procedures which carry out symbolic and execution time preprocessing needed for data exchange in block structured and structured adaptive programs were developed last summer and are being refined by Alan Sussman and S. Gupta. Primitives are being used to implement and extend Fortran D in joint work with Ken Kennedy and Geoffrey Fox's groups at the NSF Center for Research on Parallel Computing. Work is also underway to use these primitives to port a version of a block structured code developed by Veer Vatsa at NASA Langley.

We have developed a scheme for carrying out highly irregular computations in which preprocessing is not possible or not profitable. A program may carry out a single pass through a concurrent, irregular loop. In this case, it is not likely to be profitable to carefully partition data and loop iterations. In addition, if the loop iterations and data have not already been carefully mapped, there may be a large number of off-processor data accesses. Memory constraints may make it impractical to store the list of off-processor data accesses associated with a loop. It may therefore be necessary to dynamically fetch and then overwrite consecutively accessed sets of off-processor data. Dynamic methods may also be called for when parallelizing loops with loop carried dependencies determined by runtime data because preprocessing may be impossible or impractical.

Our scheme contains these components:

A low latency, flexible communications interface capable of fetching, storing and modifying irregular patterns of off-processor data. The interface needs to contain control constructs so that off-processor data can be read or written when processors have reached the appropriate points in their computations.

Separate threads responsible for computation (computation threads) and for determining which off-processor data to prefetch (prefetch threads). The prefetch threads make use of

the communications interface.

Compiler transformations to produce code which generates and controls the computation and data prefetch threads.

The compiler also modifies data structures to allocate space where data fetch threads will place off-processor data copies. In addition, the compiler also modifies loop structure so that on-processor data and off-processor copies are properly accessed. We are also considering hardware support to facilitate access to copies of off-processor data. Our schemes will function without the hardware support but will be accelerated by our proposed hardware supported mechanisms for rapidly addressing off-processor data copy storage.

We have had initial success in developing the infrastructure required for the developing a low latency, flexible communications interface, the other two components our scheme remain to be developed. This work was done in cooperation with Matt Rosing.

Alan Sussman

We have been working on the development of tools for mapping multi-block multigrid applications onto distributed memory parallel processors. With Satyanarayan Gupta, I have implemented a library of routines that provides Fortran D style declarations of distributed arrays, for both Fortran and C programs, along with routines for managing the communication required to maintain globally consistent distributed arrays. The library is based on Kay Crowley and Craig Chase's initial implementation. We are incorporating these routines into a multi-block, multigrid Navier Stokes code written by V. Vatsa and M. Sanetrik (Fluid Mechanics Division, LaRC), with the eventual goal of building a compiler that can transform a sequential block structured program into a distributed memory parallel program.

We have also begun to investigate the application of the block structured library routines to structured adaptive problems. We will be working with a structured adaptive grid program developed by J. Quirk, to augment our current set of routines with the additional functionality required to support adaptive grids on a distributed memory parallel machine.

I have also been working with T. Eidson on investigating the behavior of floating point arithmetic on the nodes of the iPSC/860 hypercube. We are using Kahan's Paranoia program to help determine the differences between strict IEEE conformance and the (faster) non-IEEE i860 option for various floating point operations.

John Van Rosendale

While mapping of standard multigrid to highly parallel architectures has been well studied, new issues arise in mapping more complex multigrid algorithms, based on line- or

plane-relaxation or based on multiple coarse grids. These newer algorithms are essential to achieving fast convergence for problems having stretched grids or anisotropic operators. Thus parallel implementation of these algorithm is important in many applications areas including computational fluid dynamics.

In collaboration with Dave Nicol and Andrea Overman, the author is looking at the mapping issues involved in two versions of the MSG (multiple semicoarsened grid) multigrid algorithm. One version uses the standard multigrid V-cycle, in which the smoothing iteration is done on only one multigrid level at a time. This is an effective numerical algorithm, but raises difficult mapping issues on distributed memory architectures. There is also a more parallel version of this algorithm, using a C-cycle in which all levels are smoothed simultaneously. This version greatly simplifies these complex mapping issues but is numerically less efficient.

One approach being explored to this mapping problem in the difficult V-cycle case, is use of dynamic programming. In MSG, one has multiple grids on each level, each having a different shape. It can be shown that dynamic programming assigns mesh points to processors optimally, assuming that each processor is used for only one grid on a level, an assumption which is reasonable for parallel machines having hundreds or thousands of processors. Once we have finished this theoretical work on mapping, we plan to implement the resulting algorithms on the Intel Delta or Paragon or on a CM-5.

Mohammed Zubair

Together with M. Ghose (Old Dominion University) we have completed a preliminary study on the performance of sparse Cholesky factorization on INTEL iPSC/860. The problem of Cholesky factorization of a sparse matrix has been very well investigated on sequential machines. A number of efficient codes exist for factorizing large unstructured sparse matrices, for example, codes from Harwell Subroutine Library and Sparspak. However, there is a lack of such efficient codes on parallel machines in general, and distributed memory machines in particular. Some of the issues which are critical to the implementation of sparse Cholesky factorization on a distributed memory parallel machine are: ordering, partitioning and mapping, load balancing, and ordering of various tasks within a processor. Addressing these issues optimally for unstructured sparse matrices is a challenging task. In this study we focused on the effect of various partitioning schemes on the performance of sparse Cholesky factorization on the INTEL iPSC/860. We also proposed a new partitioning heuristic for structured as well as unstructured sparse matrices, and compare its performance with the other schemes.

REPORTS AND ABSTRACTS

Jerome, Joseph W., and Chi-Wang Shu: *Energy models for one-carrier transport in semiconductor devices*. ICASE Report No. 91-78, October 10, 28 pages. To appear in IMA Volume in Mathematics and Its Applications.

Moment models of carrier transport, derived from the Boltzmann equation, have made possible the simulation of certain key effects through such realistic assumptions as energy dependent mobility functions. This type of global dependence permits the observation of velocity overshoot in the vicinity of device junctions, not discerned via classical drift-diffusion models, which are primarily local in nature. It has been found that a critical role is played in the hydrodynamic model by the heat conduction term. When ignored, the overshoot is inappropriately damped. When the standard choice of the Wiedemann-Franz law is made for the conductivity, spurious overshoot is observed. Agreement with Monte-Carlo simulation in this regime has required empirical modification of this law, as observed by IBM researchers, or nonstandard choices. In this paper, simulations of the hydrodynamic model in one and two dimensions, as well as simulations of a newly developed energy model, the RT model, will be presented. The RT model, intermediate between the hydrodynamic and drift-diffusion model, was developed at the University of Illinois to eliminate the parabolic energy band and Maxwellian distribution assumptions, and to reduce the spurious overshoot with physically consistent assumptions. The algorithms employed for both models are the essentially non-oscillatory shock capturing algorithms, developed at UCLA during the last decade. Some mathematical results will be presented, and contrasted with the highly developed state of the drift-diffusion model.

Nicol, David, Rahul Simha, Alok Choudhury and Bhagirath Narahari: *Optimal processor assignment for pipeline computations*. ICASE Report No. 91-79, October 10, 1991, 36 pages. Submitted to IEEE Trans. on Parallel & Distributed Systems.

The availability of large scale multitasked parallel architectures introduces the following processor assignment problem for pipelined computations. Given a set of tasks and their precedence constraints, along with their experimentally determined individual response times for different processor sizes, find an assignment of processors to tasks. Two objectives interest us: minimal response given a throughput requirement, and maximal throughput given a response time requirement. These assignment problems differ considerably from the classical mapping problem in which several tasks share a processor; instead, we assume that a large number of processors are to be assigned to a relatively small number of tasks. In this paper we develop efficient assignment algorithms for different classes of task structures. For a p processor system and a series-parallel precedence graph with n constituent tasks, we provide an $O(np^2)$ algorithm that finds the optimal assignment for the response time optimization problem; we find the assignment optimizing the constrained throughput in $O(np^2 \log p)$ time. Special cases of linear, independent, and tree graphs are also considered. In addition, we also examine more efficient algorithms when certain restrictions are placed on the problem parameters. Our techniques are applied to a task system in computer vision.

Venugopal, Sesh, and Vijay K. Naik: *Effects of partitioning and scheduling sparse matrix factorization on Communication and load balance*. ICASE Report No. 91-80, October 10, 1991, 20 pages. To appear in Proceedings of Supercomputing 1991.

We present a block-based, automatic partitioning and scheduling methodology for sparse matrix factorization on distributed memory systems. Using experimental results, we analyze this technique for communication and load imbalance overhead. To study the performance effects, we compare these overheads with those obtained from a straightforward “wrap-mapped” column assignment scheme. All experimental results were obtained using test sparse matrices from the Harwell-Boeing data set. The results show that there is a communication and load balance trade-off. The block-based method results in lower communication cost whereas the wrap-mapped scheme gives better load balance.

Bassom, Andrew P., and Sharon O. Seddougui: *The effects of suction on the nonlinear stability of the three-dimensional boundary layer above a rotating disc*. ICASE Report No. 91-81, November 5, 1991, 16 pages. Submitted to Proceedings of the Royal Society of London, Series A.

There exist two types of stationary instability of the flow over a rotating disc corresponding to the upper-branch, inviscid mode and the lower-branch mode, which has a triple deck structure, of the neutral stability curve. A theoretical investigation of the linear problem and an account of the weakly nonlinear properties of the lower branch-modes have been undertaken by Hall (1986) and MacKerrell (1987) respectively. Motivated by recent reports of experimental sightings of the lower-branch mode and an examination of the role of suction on the linear stability properties of the flow here we investigate the effects of suction on the nonlinear disturbance described by MacKerrell (1987). The additional analysis required in order to incorporate suction is relatively straightforward and enables us to derive an amplitude equation which describes the evolution of the mode. For each value of the suction a threshold value of the disturbance amplitude is obtained; modes of size greater than this threshold grow without limit as they develop away from the point of neutral stability.

Zhang, H.S., R.M.C. So, C.G. Speziale and L.G. Lai: *A near-wall two-equation model for compressible turbulent flows*. ICASE Report No. 91-82, November 4, 1991, 37 pages. To be submitted to AIAA Journal.

A near-wall two-equation turbulence model of the $K - \varepsilon$ type is developed for the description of high-speed compressible flows. The Favre-averaged equations of motion are solved in conjunction with modeled transport equations for the turbulent kinetic energy and solenoidal dissipation wherein a variable density extension of the asymptotically consistent near-wall model of So and co-workers is supplemented with new dilatational models. The resulting compressible two-equation model is tested in the supersonic flat plate boundary layer – with an adiabatic wall and with wall cooling – for Mach numbers as large as 10. Direct comparisons of the predictions of the new model with raw experimental data and with results from the $K - \omega$ model indicate that it performs well for a wide range of Mach numbers. The surprising finding is that the Morkovin hypothesis, where turbulent dilatational terms are neglected, works well at high Mach numbers provided that the near wall model is asymptotically consistent. Instances where the model predictions deviate from the experiments appear

to be attributable to the assumption of constant turbulent Prandtl number – a deficiency that will be addressed in a future paper.

Nicol, David M., Albert G. Greenberg and Boris D. Lubachevsky: *Massively parallel algorithms for trace-driven cache simulations*. ICASE Report No. 91-83, November 5, 1991, 22 pages. Submitted to IEEE Trans. on Parallel and Distributed Systems (journal).

Trace-driven cache simulation is central to computer design. A trace is a very long sequence, x_1, \dots, x_n , of references to lines (contiguous locations) from main memory. At the t^{th} instant, reference x_t is hashed into a set of cache locations, the contents of which are then compared with x_t . If at the t^{th} instant x_t is not present in the cache, then it is said to be a *miss*, and is loaded into the cache set, possibly forcing the replacement of some other memory line, and making x_t present for the $(t + 1)^{\text{st}}$ instant. The problem of parallel simulation of a subtrace of N references are misses and related statistics.

A simulation method is presented for the Least-Recently-Used (LRU) policy, which regardless of the set size C runs in time $O(\log N)$ using N processors on the exclusive read, exclusive write (EREW) parallel model. A simpler LRU simulation algorithm is given that runs in $O(C \log N)$ time using $N/\log N$ processors. We present timings of the second algorithm's implementation on the MasPar MP-1, a machine with 16384 processors. A broad class of *reference-based* line replacement policies are considered, which includes LRU as well as the Least-Frequently-Used and Random replacement policies. A simulation method is presented for any such policy that on any trace of length N directed to a C line set runs in time $O(C \log N)$ time with high probability using N processors on the EREW model. The algorithms are simple, have very little space overhead, and are well-suited for SIMD implementation.

Abgrall, Remi: *Design of an essentially non-oscillatory reconstruction procedure on finite-element type meshes*. ICASE Report No. 91-84, December 10, 1991, 34 pages. Submitted to the 8th International Conference on Finite Element Methods, Swansea, March 1992.

In this report, we have designed an essentially non-oscillatory reconstruction for functions defined on finite-element type meshes. Two related problems are studied : the interpolation of possibly unsmooth multivariate functions on arbitrary meshes and the reconstruction of a function from its average in the control volumes surrounding the nodes of the mesh. Concerning the first problem, we have studied the behaviour of the highest coefficients of the Lagrange interpolation function which may admit discontinuities of locally regular curves. This enables us to choose the best stencil for the interpolation. The choice of the smallest possible number of stencils is addressed. Concerning the reconstruction problem, because of the very nature of the mesh, the only method that may work is the so called reconstruction via deconvolution method. Unfortunately, it is well suited only for regular meshes as we show, but we also show how to overcome this difficulty. The global method has the expected order of accuracy but is conservative up to a high order quadrature formula only.

Some numerical examples are given which demonstrate the efficiency of the method.

Morris, K.A.: *Convergence of controllers designed using state space methods*. ICASE Report No. 91-85, December 10, 1991, 20 pages. Submitted to SIAM Journal on Control & Optimization.

In this paper convergence of finite-dimensional controllers for infinite-dimensional systems designed using approximations is examined. Stable coprime factorization theory is used to show that under the standard assumptions of uniform stabilizability/detectability, the controllers stabilize the original system for large enough model order. The controllers converge uniformly to an infinite-dimensional controller, as does the closed loop response.

Denier, James P., and Philip Hall: *On the nonlinear development of the most unstable Görtler vortex mode*. ICASE Report No. 91-86, December 10, 1991, 33 pages. Submitted to J. Fluid Mechanics.

The nonlinear development of the most unstable Gvortex mode in boundary layer flows over curved walls is investigated. The most unstable Gmode is confined to a viscous wall layer of thickness $O(G^{-1/5})$ and has spanwise wavelength $O(G^{-1/5})$; it is, of course, most relevant to flow situations where the Gnumber $G \gg 1$. The nonlinear equations governing the evolution of this mode over an $O(G^{-3/5})$ streamwise lengthscale are derived and are found to be of a fully nonparallel nature. The solution of these equations is achieved by making use of the numerical scheme used by Hall (1988) for the numerical solution of the nonlinear Gequations valid for $O(1)$ Gnumbers. Thus, the spanwise dependence of the flow is described by a Fourier expansion whereas the streamwise and normal variations of the flow are dealt with by employing a suitable finite difference discretization of the governing equations. Our calculations demonstrate that, given a suitable initial disturbance, after a brief interval of decay, the energy in all the higher harmonics grows until a singularity is encountered at some downstream position. The structure of the flow field as this singularity is approached suggests that the singularity is responsible for the vortices, which are initially confined to the thin viscous wall layer, moving away from the wall and into the core of the boundary layer.

Fu, Yibin, and Philip Hall: *Effects of Görtler vortices, wall cooling and gas dissociation on the Rayleigh instability in a hypersonic boundary layer*. ICASE Report No. 91-87, December 11, 1991, 33 pages. Submitted to J. Fluid Mechanics.

In a hypersonic boundary layer over a wall of variable curvature, the region most susceptible to Görtler vortices is the temperature adjustment layer sitting at the edge of the boundary layer (Hall & Fu (1989), Fu, Hall & Blackaby (1990)). This temperature adjustment layer is also the most dangerous site for Rayleigh instability (Cowley & Hall (1990), Smith & Brown (1990) and Blackaby, Cowley and Hall (1990)). In this paper, we investigate how the existence of large amplitude Görtler vortices affects the growth rate of Rayleigh instability. The effects of wall cooling and gas dissociation on this instability are also studied. We find that all these mechanisms increase the growth rate of Rayleigh instability and are therefore destabilizing.

Banks, H.T., W. Fang, R.J. Silcox and R.C. Smith: *Approximation methods for control of acoustic/structure models with piezoceramic actuators*. ICASE Report No. 91-88, December 11, 1991, 30 pages. Submitted to Journal of Intelligent Material Systems and Structures.

The active control of acoustic pressure in a 2-D cavity with a flexible boundary (a beam) is considered. Specifically, this control is implemented via piezoceramic patches on the beam which produce pure bending moments. The incorporation of the feedback control in this manner leads to a system with an unbounded input term. Approximation methods in the context of an LQR state space formulation are discussed and numerical results demonstrating the effectiveness of this approach in computing feedback controls for noise reduction are presented.

Radespiel, Rolf, and R.C. Swanson: *Progress with multigrid schemes for hypersonic flow problems*. ICASE Report No. 91-89, December 16, 1991, 42 pages. To be submitted to Journal of Computational Physics.

Several multigrid schemes are considered for the numerical computation of viscous hypersonic flows. For each scheme, the basic solution algorithm employs upwind spatial discretization with explicit multistage time stepping. Two-level versions of the various multigrid algorithms are applied to the two-dimensional advection equation, and Fourier analysis is used to determine their damping properties. The capabilities of the multigrid methods are assessed by solving three different hypersonic flow problems. Some new multigrid schemes based on semicoarsening strategies are shown to be quite effective in relieving the stiffness caused by the high-aspect-ratio cells required to resolve high Reynolds number flows. These schemes exhibit good convergence rates for Reynolds numbers up to 200×10^6 and Mach numbers up to 25.

1992

Chase, Craig, Kay Crowley, Joel Saltz and Anthony Reeves: *Parallelization of irregularly coupled regular meshes*. ICASE Report No. 92-1, January 7, 1992, 25 pages. International Supercomputing Conference.

Regular meshes are frequently used for modeling physical phenomena on both serial and parallel computers. One advantage of regular meshes is that efficient discretization schemes can be implemented in a straightforward manner. However, geometrically-complex objects, such as aircraft, cannot be easily described using a single regular mesh. Multiple interacting regular meshes are frequently used to describe complex geometries. Each mesh models a subregion of the physical domain. The meshes, or *subdomains*, can be processed in parallel, with periodic updates carried out to move information between the coupled meshes. In many cases, there are a relatively small number (one to a few dozen) subdomains, so that each subdomain may also be partitioned among several processors.

We outline a composite run-time/compile-time approach for supporting these problems efficiently on distributed-memory machines. This paper describes these methods in the

context of a multiblock fluid dynamics problem developed at the NASA Langley Research Center.

E, Weinan, and Chi-Wang Shu: *Effective equations and the inverse cascade theory for Kolmogorov flows*. ICASE Report No. 92-2, January 7, 1992, 22 pages. To be submitted to Physics of Fluids A.

We study the two dimensional Kolmogorov flows in the limit as the forcing frequency goes to infinity. Direct numerical simulation indicates that in this limit the low frequency energy spectrum evolves to a universal k^{-4} decay law. We derive effective equations governing the behavior of the large scale flow quantities. We then present numerical evidence that with smooth initial data, the solution to the effective equation develops a k^{-4} type singularity at a finite time. This gives a convenient explanation for the k^{-4} decay law exhibited by the original Kolmogorov flows.

Speziale, C. G., and S. Thangam: *Analysis of an RNG based turbulence model for separated flows*. ICASE Report No. 92-3, January 23, 1992, 18 pages. To appear in the International Journal of Engineering Science.

A two-equation turbulence model of the $K - \epsilon$ type was recently derived by Yakhot & Orszag based on Renormalization Group (RNG) methods. It was later reported that this RNG based model yields substantially better predictions than the standard $K - \epsilon$ model for turbulent flow over a backward facing step – a standard test case used to benchmark the performance of turbulence models in separated flows. The apparent improvements obtained from the RNG $K - \epsilon$ model were attributed to the better treatment of near wall turbulence effects. In contrast to these earlier claims, it is shown in this paper that the original version of the RNG $K - \epsilon$ model substantially underpredicts the reattachment point in the backstep problem – a deficiency that is traced to the modeling of the production of dissipation term. However, with the most recent improvements in the RNG $K - \epsilon$ model proposed by Yakhot and co-workers, excellent results for the backstep problem are now obtained. Interestingly enough, these results are not that sensitive to the details of the near wall treatment.

Gottlieb, David, Chi-Wang Shu, Alex Solomonoff, and Herve Vandeven: *On the Gibbs phenomenon I: Recovering exponential accuracy from the fourier partial sum of a non-periodic analytic function*. ICASE Report No. 92-4, February 27, 26 pages. Submitted to Journal of Computational and Applied Mathematics.

It is well known that the Fourier series of an analytic and periodic function, truncated after $2N + 1$ terms, converges *exponentially* with N , even in the maximum norm. It is also known that if the function is *not* periodic, the rate of convergence deteriorates; in particular there is no convergence in the maximum norm, although the function is still analytic. This is known as the **Gibbs phenomenon**. In this paper we show that the first $2N + 1$ Fourier coefficients contain enough information about the function so that an exponentially convergent approximation (in the *maximum* norm) can be constructed. The

proof is a constructive one and makes use of the Gegenbauer polynomials $C_n^\lambda(x)$. It consists of the following two steps.

In the first step we show that the first m coefficients of the Gegenbauer expansion (based on $C_n^\lambda(x)$, for $0 \leq n \leq m$) of any L_2 function can be obtained, within exponential accuracy, provided that both λ and m are proportional to (but smaller than) N .

In the second step we construct the Gegenbauer expansion based on C_n^λ , $0 \leq n \leq m$ from the coefficients found in the first step. We show that this series converges exponentially with N , provided that the original function is analytic (though non-periodic).

Thus we prove that : **The Gibbs phenomenon can be completely overcome.**

Lafon, F., and S. Osher: *Essentially nonoscillatory postprocessing filtering methods*. ICASE Report No. 92-5, January 31, 1992, 23 pages. To appear in Proceedings of Algorithmic Trends for Computational Fluid Dynamics in the 90's.

High order accurate centered flux approximations used in the computation of numerical solutions to nonlinear partial differential equations produce large oscillations in regions of sharp transitions. In this paper, we present a new class of filtering methods denoted by ENO-LS (Essentially Nonoscillatory Least Squares) which constructs an upgraded filtered solution that is close to the physically correct weak solution of the original evolution equation. Our method relies on the evaluation of a least squares polynomial approximation to oscillatory data using a set of points which is determined via the ENO framework.

Numerical results are given in one and two space dimensions for both scalar and systems of hyperbolic conservation laws. Computational running time, efficiency and robustness of the method are illustrated in various examples such as Riemann initial data for both Burgers' and Euler's equations of gas dynamics. In all standard cases the filtered solution appears to converge numerically to the correct solution of the original problem. Some interesting results based on nonstandard central difference schemes, which exactly preserve entropy, and have been recently shown generally not to be weakly convergent to a solution of the conservation law, are also obtained using our filters.

Sarkar, S., G. Erlebacher, and M. Y. Hussaini: *Compressible homogeneous shear: Simulation and modeling*. ICASE Report No. 92-6, February 10, 1992, 27 pages. To appear in Turbulent Shear Flows 8: Selected Papers.

The present study investigates compressibility effects on turbulence by direct numerical simulation of homogeneous shear flow. A primary observation is that the growth of the turbulent kinetic energy decreases with increasing turbulent Mach number. The sinks provided by compressible dissipation and the pressure-dilatation, along with reduced Reynolds shear stress, are shown to contribute to the reduced growth of kinetic energy. Models are proposed for these dilatational terms and verified by direct comparison with the simulations. The differences between the incompressible and compressible fields are brought out by the examination of spectra, statistical moments, and structure of the rate of strain tensor.

Quirk, James J.: *An alternative to unstructured grids for computing gas dynamic flows around arbitrarily complex two-dimensional bodies*. ICASE Report No. 92-7, February 27, 1992, 31 pages. Submitted to Computers and Fluids.

Within the shock-capturing community, the need to simulate flows around geometrically complex bodies has resulted in an inexorable shift away from schemes which employ body-fitted grids to schemes which employ unstructured grids. Although unstructured grids are undeniably effective, in view of the increasing reliance placed on computational results, such a wholesale shift in mentality should give cause for concern. The concept of using several computer codes to cross check numerical results becomes ill-founded if all codes follow the same methodology. In this paper we describe an alternative approach for dealing with arbitrarily complex, two-dimensional geometries, the so-called cartesian boundary method.

Conceptually, the cartesian boundary method is quite simple. Solid bodies blank out areas of a background, cartesian mesh, and the resultant cut cells are singled out for special attention. However, there are several obstacles that must be overcome in order to achieve a practical scheme. We present a general strategy that overcomes these obstacles, together with some details of our successful conversion of an adaptive mesh algorithm from a body-fitted code to a cartesian boundary code.

Sussman, Alan: *Execution models for mapping programs onto distributed memory parallel computers*. ICASE Report No. 92-8, March 3, 1992, 30 pages. Submitted to the 20th International Conference on Parallel Processing – August 1992.

This paper addresses the problem of exploiting the parallelism available in a program to efficiently employ the resources of the target machine, in the context of building a mapping compiler for a distributed memory parallel machine. The paper describes using *execution models* to drive the process of mapping a program in the most efficient way onto a particular machine.

Through analysis of the execution models for several mapping techniques for one class of programs, we show that the selection of the best technique for a particular program instance can make a significant difference in performance. On the other hand, the results of benchmarks from an implementation of a mapping compiler show that our execution models are accurate enough to select the best mapping technique for a given program.

Chapman, Barbara, Piyush Mehrotra, and Hans Zima: *Programming in Vienna Fortran*. ICASE Report No. 92-9, March 18, 1992, 40 pages. Submitted to Scientific Programming.

Exploiting the full performance potential of distributed memory machines requires a careful distribution of data across the processors. Vienna Fortran is a language extension of Fortran which provides the user with a wide range of facilities for such mapping of data structures. In contrast to current programming practice, programs in Vienna Fortran are written using global data references. Thus, the user has the advantages of a shared memory programming paradigm while explicitly controlling the data distribution. In this paper, we present the language features of Vienna Fortran for FORTRAN 77, together with examples illustrating the use of these features.

Jacobs, P. A.: *Numerical simulation of transient hypervelocity flow in an expansion tube*. ICASE Report No. 92-10, March 11, 1992, 48 pages. Submitted to the International Journal of Computers and Fluids.

Several numerical simulations of the transient flow of helium in an expansion tube are presented in an effort to identify some of the basic mechanisms which cause the noisy test flows seen in experiments. The calculations were performed with an axisymmetric Navier-Stokes code based on a finite-volume formulation and upwinding techniques. Although laminar flow and ideal bursting of the diaphragms was assumed, the simulations showed some of the important features seen in the experiments. In particular, the discontinuity in tube diameter at the primary diaphragm station introduced a transverse perturbation to the expanding driver gas and this perturbation was seen to propagate into the test gas under some flow conditions. The disturbances seen in the test flow can be characterized as either small-amplitude, low-frequency noise possibly introduced during shock compression or large-amplitude, high-frequency noise associated with the passage of the reflected head of the unsteady expansion.

Lasseigne, D.G., and M.Y. Hussaini: *Interaction of disturbances with an oblique detonation wave attached to a wedge*. ICASE Report No. 92-11, March 19, 1991, 31 pages. Submitted to Physics of Fluids A.

The linear response of an oblique overdriven detonation to imposed free stream disturbances or to periodic movements of the wedge is examined. The freestream disturbances are assumed to be steady vorticity waves and the wedge motions are considered to be time periodic oscillations either about a fixed pivot point or along the plane of symmetry of the wedge aligned with the incoming stream. The detonation is considered to be a region of infinitesimal thickness in which a finite amount of heat is released. The response to the imposed disturbances is a function of the Mach number of the incoming flow, the wedge angle, and the exothermicity of the reaction within the detonation. It is shown that as the degree of overdrive increases, the amplitude of the response increases significantly; furthermore, a fundamental difference in the dependence of the response on the parameters of the problem is found between the response to a free stream disturbance and to a disturbance emanating from the wedge surface.

Das, R., D.J. Mavriplis, J. Saltz, S. Gupta, and R. Ponnusamy: *The design and implementation of a parallel unstructured Euler solver using software primitives*. ICASE Report No. 92-12, March 18, 1992, 26 pages. Proceedings of the AIAA Reno Conference - 1992.

This paper is concerned with the implementation of a three-dimensional unstructured-grid Euler-solver on massively parallel distributed-memory computer architectures. The goal is to minimize solution time by achieving high computational rates with a numerically efficient algorithm. An unstructured multigrid algorithm with an edge-based data-structure has been adopted, and a number of optimizations have been devised and implemented in order to accelerate the parallel computational rates. The implementation is carried out by creating a set of software tools, which provide an interface between the parallelization issues and the sequential code, while providing a basis for future automatic run-time compilation support. Large practical unstructured grid problems are solved on the Intel iPSC/860 hypercube and

Intel Touchstone Delta machine. The quantitative effect of the various optimizations are demonstrated, and we show that the combined effect of these optimizations leads to roughly a factor of three performance improvement. The overall solution efficiency is compared with that obtained on the CRAY-YMP vector supercomputer.

Zubair, M., and M. Ghose: *A performance study of sparse Cholesky factorization on Intel iPSC/860*. ICASE Report No. 92-13, March 25, 1992, 18 pages. Submitted to Frontiers '92: The 4th Symposium on the Frontiers of Massively Parallel Computation; McLean, Virginia.

The problem of Cholesky factorization of a sparse matrix has been very well investigated on sequential machines. A number of efficient codes exist for factorizing large unstructured sparse matrices, for example, codes from Harwell Subroutine Library [?] and Sparspak [?]. However, there is a lack of such efficient codes on parallel machines in general, and distributed memory machines in particular. Some of the issues which are critical to the implementation of sparse Cholesky factorization on a distributed memory parallel machine are: *ordering, partitioning and mapping, load balancing, and ordering of various tasks* within a processor. Addressing these issues optimally for unstructured sparse matrices is a challenging task. In this paper we focus on the effect of various partitioning schemes on the performance of sparse Cholesky factorization on the INTEL iPSC/860. We also propose a new partitioning heuristic for structured as well as unstructured sparse matrices, and compare its performance with the other schemes.

ICASE INTERIM REPORTS

Jacobs, P.A.: *Numerical simulation of transient hypervelocity flow in a n expansion tube.*, Interim Report No. 20, January 22, 1992, 57 pages.

We present several numerical simulations of the transient flow of helium in an expansion tube. The aim of the exercise is to provide further information on (and hopefully further insight into) the operational problems of the NASA Langley expansion tube. The calculations were performed with an axi symmetric Navier-Stokes code based on a finite-volume formulation and upwinding techniques. Although laminar flow and ideal bursting of the diaphragms was assumed, the simulations showed some of the important features seen in the experiments. In particular, the discontinuity in tube diameter at the primary diaphragm station introduced a transverse perturbation to the expanding driver gas and this perturbation was seen to propagate into the test gas under some flow conditions. The disturbances seen in the test flow can be characterized as either "small-amplitude" noise possibly introduced during shock compression or "large-amplitude" noise associated with the passage of the reflected head of the unsteady expansion.

Zima, Hans, Peter Brezany, Barbara Chapman, Piyush Mehrotra, and Andreas Schwald: *Vienna Fortran - A language specification version 1.1.* Interim Report No. 21, March 19, 1992, 89 pages.

This document presents the syntax and semantics of Vienna Fortran, a machine-independent language extension to FORTRAN 77, which allows the user to write programs for distributed-memory systems using global addresses. Vienna Fortran includes high-level features for specifying virtual processor structures, distributing data across sets of processors, dynamically modifying distributions, and formulating explicitly parallel loops. The language is based upon the Single-Program-Multiple-Data (SPMD) paradigm, which exploits the parallelism inherent in many scientific codes. A substantial subset of the language features has already been implemented.

ICASE COLLOQUIA

April 1, 1991 - September 30, 1991

Name/Affiliation/Title	Date
T.C. Corke, Illinois Institute of Technology "Resonance in Axisymmetric Jets with Controlled Helical Mode Input"	October 10
Moshe Matalon, Northwestern University "Propagation, Extinction and Stability of Stretched Flames"	October 30
Peter Protzel, ICASE "Adaptive Neural Network Controller for Fault-Tolerant Autonomous Systems"	November 13
Freda Porter-Locklear, Pembroke State University "Error Propagation into C^∞ Regions"	December 6
A. Leonard, California Institute of Technology "Study of Three-Dimensional Transitions in Spherical Couette Flow Using Numerical Simulation"	December 9
John F. Clarke, Cranfield Institute of Technology "High-Speed Combustion"	December 11
Andreas Acrivos, The Levich Institute, City College of CUNY "The Stability in a 2-D Hagen-Poiseuille Resuspension Flow"	December 17
J.P. Bonnet, Centre D'Etudes Aerodynamiques et Thermiques, France "Some Experimental Evidences of Compressibility Effect in Supersonic Turbulent Wakes and Conventional Second-Order Closure Performances"	January 22
Jackson Herring, National Center for Atmospheric Sciences "An Assessment of Vortex Structures on Turbulent Spectra"	January 23
S.P. Lin, Clarkson University "Linear and Nonlinear Stability of Two Exact Periodic Nonparallel Flows"	February 14

Name/Affiliation/Title	Date
Abdon Sepulveda, University of California, Los Angeles "On the Solution of Problems with Nonconvexities and Discrete Variables in Engineering Optimization"	February 21
Hans P. Zima and Barbara Chapman, University of Vienna, Austria "Programming in Vienna Fortran"	February 27
Stephen Orszag, Princeton University "Synergistic Interactions between Theory and Computations in Fluid Dynamics"	March 6
Leslie Smith, Princeton University "Fundamentals and Applications of the Renormalization Group Theory of Turbulence"	March 13
Scott T. Leutenegger, IBM T.J. Watson Research Center "Processor Allocation for Multiprogrammed Multiprocessor Scheduling"	March 19
Terrence W. Pratt, University of Virginia "A New Solution to the 'Dusty Deck' Problem for MIMD Parallel Computers"	March 20
Ravi Jain, University of Texas at Austin "Scheduling Data Transfers in Parallel Computing Systems"	March 24
Robert P. Weaver, University of Colorado at Boulder "Providing Support for Unstructured and/or Dynamic Data in Languages for Distributed Memory Machines"	March 25
Ronald Rehm, National Institute of Standards and Technology "Finite-Rate Diffusion-Controlled Reaction in a Vortex"	March 30

ICASE STAFF

I. ADMINISTRATIVE

M. Yousuff Hussaini, Acting Director, beginning November 1, 1991 through April 2, 1992. (Director, beginning April 3, 1992) Ph.D., Mechanical Engineering, University of California, 1970 Computational Fluid Dynamics.

Robert G. Voigt, Director, through October 10, 1991. Ph.D., Mathematics, University of Maryland, 1969. Numerical and Algorithms for Parallel Computers

Linda T. Johnson, Office and Financial Administrator

Etta M. Blair, Personnel/Bookkeeping Secretary

Barbara A. Cardasis, Administrative Secretary

Tamiko J. Hackett, Office Clerk

Rachel A. Lomas, Bookkeeping Clerk

Rosa H. Milby, Short-term Housing/Office Secretary

Shelly D. Millen, Technical Publications Secretary

Emily N. Todd, Executive Secretary/Visitor Coordinator

II. ASSOCIATE MEMBERS

Saul S. Abarbanel, Professor, Department of Applied Mathematics, Tel-Aviv University, Israel.

H. Thomas Banks, Professor, Center for Applied Mathematical Sciences, University of Southern California.

David Gottlieb, Professor, Division of Applied Mathematics, Brown University.

Peter D. Lax, Professor, Courant Institute of Mathematical Sciences, New York University.

III. CHIEF SCIENTIST

M. Yousuff Hussaini - Ph.D., Mechanical Engineering, University of California, 1970. Computational Fluid Dynamics. (Beginning April 1978)

IV. LEAD COMPUTER SCIENTIST

Joel H. Saltz - Ph.D., Computer Science, Duke University, 1985. Performance optimization, programming tools and compilers for irregular and adaptive problems on scalable multiprocessors. (Beginning July 1989)

V. SENIOR STAFF SCIENTIST

Gordon Erlebacher - Ph.D., Plasma Physics, Columbia University, 1983. Computational Fluid Dynamics. (November 1989 to November 1994)

Dimitri Mavriplis - Ph.D., Mechanical and Aerospace Engineering, Princeton University, 1988. Grid Techniques for Computational Fluid Dynamics. (February 1987 to September 1995)

Piyush Mehrotra - Ph.D., Computer Science, University of Virginia, 1982. Programming Languages and Compilers for Multiprocessor Systems. (January 1991 to September 1992)

Charles G. Speziale - Ph.D., Mechanical and Aerospace Engineering, Princeton University, 1978. Fluid Dynamics with Emphasis on Turbulence Modeling and the Transition Process. (September 1987 to September 1992)

Shlomo Ta'asan - Ph.D., Applied Mathematics, The Weizmann Institute of Science, 1985. Multigrid Methods for Partial Differential Equations. (July 1991 to July 1994)

John R. Van Rosendale - Ph.D., Computer Science, University of Illinois, 1980. Parallel Systems and Algorithms. (July 1989 to July 1992)

VI. SCIENTIFIC STAFF

Fabio Bertolotti - Ph.D., Mechanical Engineering, Ohio State University, 1991. Stability Theory in Fluid Mechanics. (September 1991 to September 1993)

Kurt M. Bryan - Ph.D., Mathematics, University of Washington, 1990. Theoretical and Computational Methods for Inverse Problems. (August 1990 to September 1993)

Leon M. Clancy - B.S., Mechanical Engineering, University of Washington, 1971. System Manager. (December 1989 to Present)

Thomas W. Crockett - B.S., Mathematics, College of William and Mary, 1977. Parallel Visualization Research. (February 1987 to September 1992)

Subhendu Das - M.S., Computer Science, College of William and Mary, 1990. Parallel Tools and Environments for Unstructured Scientific Computations. (June 1991 to June 1993)

Stephen Otto - Ph.D., Mathematics, Exeter University, England, 1991. Asymptotic Theories of Stability and Transition. (November 1991 to November 1993)

Yuh-Roung Ou - Ph.D., Aerospace Engineering, University of Southern California, 1988. Control Systems for Fluid Dynamics. (November 1988 to January 1992)

Peter W. Protzel - Ph.D., Electrical Engineering, Technical University of Braunschweig, Germany, 1987. Reliability of Computing Systems. (March 1987 to November 1991)

James Quirk - Ph.D., Computational Fluid Dynamics, Cranfield Institute of Technology, 1991. Adaptive Methods for Partial Differential Equations. (June 1991 to June 1993)

Mathew Rosing - Ph.D., Computer Science, University of Colorado, Boulder, 1991. Tools and Compilers for Scalable Multiprocessors. (November 1991 to November 1994)

Sutanu Sarkar - Ph.D., Mechanical and Aerospace Engineering, Cornell University, 1988. Fluid Dynamics, Turbulence. (September 1988 to September 1992)

Ralph C. Smith - Ph.D., Numerical Analysis, Montana State University, 1990. Theoretical and Computational Methods for Feedback Control and Parameter Estimation. (August 1990 to August 1993)

Alan Sussman - Ph.D., Computer Science, Carnegie-Mellon University, 1991. Quantifying and Predicting Performance of Realistic Applications on Distributed Memory Parallel Computers. (August 1991 to August 1993)

VII. VISITING SCIENTISTS

Dave Broutman - Ph.D., Physical Oceanography, Scripps Institute of Oceanography, 1982. Lecturer/Research Fellow, Department of Mathematics, University of New South Wales, Australia. Multi-Domain Spectral Methods. (January 1992)

Barbara Chapman - M.S., Mathematics, University of Canterbury, Christchurch, New Zealand, 1985. Research Associate, Computer Science Department, University of Vienna. Compiler Development for Multiprocessors. (February 1992)

Peter W. Duck - Ph.D., Fluid Mechanics, University of Southampton, United Kingdom, 1975. Lecturer in Mathematics, Department of Mathematics, University of Manchester, United Kingdom. Stability Theory. (December 1991 to January 1992)

Marie Pierre Liandrat - Ph.D., Fluid Mechanics, University Aix-Marseille, France, 1988. Scientist, IMST, France. Wavelet Techniques to Process Numerical and Experimental Data Sets. (November to December 1991)

Merrell L. Patrick - Ph.D., Mathematics, Carnegie-Mellon University, 1964. New Technologies Program Director, National Science Foundation. (August 1991 to August 1992)

Eitan Tadmor - Ph.D., Numerical Analysis, Tel-Aviv University, 1979. Professor, Department of Applied Mathematics, Tel-Aviv University. Numerical Methods for Partial Differential Equations. (January 1992)

VIII. CONSULTANTS

Dinshaw Balsara - Ph.D., Computational Fluid Dynamics, Astro Physics, University of Illinois at Urbana, 1990. Department of Physics and Astronomy, Johns Hopkins University. Parallel Implementation of Adaptive Godunov Schemes.

Alvin Bayliss - Ph.D., Mathematics, New York University, 1975. Associate Professor, Technological Institute, Northwestern University. Numerical Methods for Partial Differential Equations.

Marsha J. Berger - Ph.D., Numerical Analysis, Stanford University, 1982. Research Associate, Courant Institute of Mathematical Sciences. Numerical Methods for Partial Differential Equations.

Percy Bobbitt - B.S., Aeronautics, Catholic University of America, 1949. NASA Langley Research Center - Retired. Aerodynamics.

John D. Buckmaster - Ph.D., Applied Mathematics, Cornell University, 1969. Professor, Department of Aeronautical and Astronautical Engineering, University of Illinois. Mathematical Combustion.

John A. Burns - Ph.D., Mathematics, University of Oklahoma, 1973. Professor, Virginia Polytechnic Institute and State University. Numerical Methods in Feedback Control and Parameter Estimation.

Jeffrey L. Clark - B.A., Business Administration, James Madison University, 1986. Programmer/Analyst, Unisys Corporation. Administration of the ICASE Computer System.

Thomas C. Corke - Ph.D., Mechanical and Aerospace Engineering, Illinois Institute of Technology, 1981. Professor, Department of Mechanical and Aerospace Engineering, Armour College of Engineering, Illinois Institute of Technology. Instability and Transition to Turbulence Experiments.

Ayodeji O. Demuren - Ph.D., Mechanical Engineering, Imperial College London, United Kingdom, 1979. Associate Professor, Department of Mechanical Engineering and Mechanics, Old Dominion University. Modeling of Turbulent Flows.

Peter R. Eiseman - Ph.D., Mathematics, University of Illinois, 1970. Senior Research Scientist and Adjunct Professor, Department of Applied Physics and of Nuclear Engineering, Columbia University. Computational Fluid Dynamics.

Robert E. Fennell - Ph.D., Mathematics, University of Iowa, 1969. Professor, Department of Mathematical Sciences, Clemson University. Control Theory for Multivariable Systems.

Dennis B. Gannon - Ph.D., Computer Science, University of Illinois, 1980. Associate Professor, Department of Computer Science, Indiana University. Parallel Computation.

James F. Geer - Ph.D., Applied Mathematics, New York University, 1967. Professor, Systems Science and Mathematical Sciences, Watson School of Engineering, Applied Science and Technology, SUNY-Binghamton. Perturbation Methods and Asymptotic Expansions of Solutions to Partial Differential Equations.

Chester E. Grosch - Ph.D., Physics - Fluid Dynamics, Stevens Institute of Technology, 1967. Professor, Department of Computer Science and Slover Professor, Department of

Oceanography, Old Dominion University. Computational Fluid Mechanics and Algorithms for Array Processor Computers.

Pratik Gupta - M.S., Computer Science, Virginia Commonwealth University, 1989. Systems Engineer, Simulation Associates Inc./Taurus Technologies Inc. Parallel Computing Systems.

Philip Hall - Ph.D., Mathematics, Imperial College, England, 1973. Professor, Department of Mathematics, University of Manchester, England. Stability Theory.

Amiram Harten - Ph.D., Mathematics, New York University, 1974. Associate Professor, Department of Mathematics, Tel-Aviv University, Israel. Numerical Analysis.

Daniel J. Inman - Ph.D., Mechanical Engineering, Michigan State University, 1980. Professor and Chair, Department of Mechanical and Aerospace Engineering, SUNY-Buffalo. Control Theory.

Thomas L. Jackson - Ph.D., Mathematics, Rensselaer Polytechnic Institute, 1985. Assistant Professor, Department of Mathematics, Old Dominion University. Numerical and Analytical Methods for Chemically Reacting Flows.

Antony Jameson - Ph.D., Magnetohydro-Dynamics, Cambridge University, England, 1963. James S. McDonnell Distinguished Professor, Department of Mechanical and Aerospace Engineering, Princeton University. Computational Aerodynamics.

Mark T. Jones - Ph.D., Computer Science, Duke University, 1990. Assistant Computer Scientist, MCS Division, Argonne National Labs. Parallel Algorithms for Numerical Linear Algebra.

Harry F. Jordan - Ph.D., Physics, University of Illinois, 1967. Professor Department of Electrical and Computer Engineering, University of Colorado at Boulder. Parallel Computation.

Ashwani K. Kapila - Ph.D., Theoretical and Applied Mechanics, Cornell University, 1975. Associate Professor, Department of Mathematical Sciences, Rensselaer Polytechnic Institute. Mathematical Combustion.

Edward J. Kerschen - Ph.D., Mechanical Engineering, Stanford University, 1978. Associate Professor, Department of Aerospace and Mechanical Engineering, University of Arizona. Receptivity and Stability Theory.

David E. Keyes - Ph.D., Applied Mathematics, Harvard University, 1984. Assistant Professor, Mechanical Engineering, Yale University. Parallelization of Numerical Procedures Appropriate for the Study of Combustion.

Heinz-Otto Kreiss - Ph.D., Mathematics, Royal Institute of Technology, Sweden 1960. Professor, Department of Applied Mathematics, California Institute of Technology. Numerical Analysis.

David G. Lasseigne - Ph.D., Applied Mathematics, Northwestern University, 1985. Assistant Professor, Department of Mathematics and Statistics, Old Dominion University. Mathematical and Numerical Combustion.

Anthony Leonard - Ph.D., Nuclear Engineering, Stanford University, 1963. Professor of Aeronautics, California Institute of Technology. Fluid Mechanics.

Robert W. MacCormack - M.S., Mathematics, Stanford University. Professor, Department of Aeronautics and Astronautics, Stanford University. Computational Aerodynamics.

Seema Mirchandaney - M.S., Computer Science, University of Massachusetts, Amherst, 1990. Research Programmer, Science and Technology Center, Rice University. Parallel Programming Environments.

Mark V. Morkovin - Ph.D., Applied Mathematics, University of Wisconsin, 1942. Professor Emeritus, Department of Mechanical and Aerospace Engineering, Illinois Institute of Technology. Instability and Transition to Turbulence.

Kirsten A. Morris - Ph.D., Electrical Engineering, University of Waterloo, 1989. Assistant Professor, Department of Applied Mathematics, University of Waterloo-Ontario, Canada. Control Theory.

Naomi H. Naik - Ph.D., Mathematics, University of Wisconsin-Madison, 1987. Visiting Assistant Professor, Department of Mathematics, Vassar College. Multi-Grid Methods.

David M. Nicol - Ph.D., Computer Science, University of Virginia, 1985. Professor, Department of Computer Science, College of William and Mary. Performance Analysis on Parallel Computing Systems.

Stanley J. Osher - Ph.D., Functional Analysis, New York University, 1966. Professor, Department of Mathematics, University of California at Los Angeles. Numerical Analysis.

Demetrius Papageorgiou - Ph.D., Mathematics, University of London, 1985. Assistant Professor, Department of Mathematics, New Jersey Institute of Technology. Stability Theory.

Ugo Piomelli - Ph.D., Mechanical Engineering, Stanford University 1987. Professor, Department of Mechanical Engineering, University of Maryland. Subgrid Scale Reynold's Stress Modelling and Large Eddy Simulation of Turbulent Flows.

Terrence W. Pratt - Ph.D., Mathematics/Computer Science, University of Texas at Austin, 1965. Professor, Department of Computer Science, University of Virginia. Parallel Computation.

Daniel A. Reed - Ph.D., Computer Science, Purdue University, 1983. Assistant Professor, Department of Computer Science, University of Illinois. Parallel Processing.

Helen L. Reed - Ph.D., Engineering Mechanics, Virginia Polytechnic Institute and State University, 1981. Associate Professor, Department of Mechanical Engineering, Arizona State University. Stability Theory.

Eli Reshotko - Ph.D., Aeronautics and Physics, California Institute of Technology, 1960. Interim Dean, Case Western Reserve University. High Speed Aerodynamics with an Emphasis on Transition, Turbulence and Combustion.

Paul E. Saylor - Ph.D., Mathematics, Rice University, 1986. Associate Professor, Computer Science Department, University of Illinois, Urbana. Iterative Solution of Linear Algebraic Equations and Algorithms for Parallel Computers.

Jeffrey S. Scroggs - Ph.D., Computer Science, University of Illinois at Urbana, 1988. Assistant Professor, Department of Mathematics, North Carolina State University. Domain Decomposition Techniques for Partial Differential Equations.

Chi-Wang Shu - Ph.D., Mathematics, University of California, Los Angeles, 1986. Associate Professor, Division of Applied Mathematics, Brown University. Numerical Partial Differential Equations.

Katepalli R. Sreenivason - Ph.D., Aeronautical Engineering, Indian Institute of Science, 1975. Professor and Chairman, Department of Mechanical Engineering, Yale University. Transition and Turbulence.

Saleh Tanveer - Ph.D., Applied Mathematics, California Institute of Technology, 1984. Professor, Department of Mathematics, Ohio State University. Applied Mathematics.

Lu Ting - Ph.D., Aeronautics, New York University, 1951. Professor, Courant Institute of Mathematical Sciences, New York University. Fluid Mechanics.

Eli Turkel - Ph.D., Applied Mathematics, New York University, 1970. Associate Professor, Department of Applied Mathematics, Tel-Aviv University, Israel. Computational Aerodynamics.

Bram van Leer - Ph.D., Theoretical Astrophysic, Leiden State University, The Netherlands, 1970. Professor, Department of Aerospace Engineering, University of Michigan. Computational Aerodynamics.

Yun Wang - M.S., Applied Mathematics, University of Southern California, 1986. Visiting Scientist, Center for Applied Mathematical Sciences, University of Southern California. Control Theory.

Hans Zima - Ph.D., Mathematics, University of Vienna, Austria, 1964. Professor, Computer Science Department, University of Vienna, Austria. Compiler Development for Parallel and Distributed Multiprocessors.

Mohammad Zubair - Ph.D., Computer Science, Indian Institute of Technology, New Delhi, India, 1987. Assistant Professor, Department of Computer Science, Old Dominion University. Performance of Unstructured Flow-Solvers on Multi Processor Machines.

IX. STUDENT ASSISTANTS

Michael Arras - Graduate Student at The College of William and Mary. (October 1988 to December 1991)

Cynthia C. Cokus - Graduate Student at The College of William and Mary. (June 1990 to December 1991)

Luis Gomes - Graduate Student at The College of William and Mary. (May 1991 to Present)

John Otten - Graduate Student at The College of William and Mary. (April 1992 to Present).

X. GRADUATE STUDENTS

Satyanarayan Gupta - Graduate Student at Old Dominion University. (June 1991 to Present)

REPORT DOCUMENTATION PAGE			Form Approved OMB No. 0704-0188	
<small>Public reporting burden for this collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden, to Washington Headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington, VA 22202-4302, and to the Office of Management and Budget, Paperwork Reduction Project (0704-0188), Washington, DC 20503.</small>				
1. AGENCY USE ONLY (Leave blank)	2. REPORT DATE May 1992	3. REPORT TYPE AND DATES COVERED Contractor Report		
4. TITLE AND SUBTITLE Semiannual Report - October 1, 1991 through March 31, 1992		5. FUNDING NUMBERS C NAS1-18605 WU 505-90-52-01		
6. AUTHOR(S)		8. PERFORMING ORGANIZATION REPORT NUMBER		
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) Institute for Computer Applications in Science and Engineering Mail Stop 132C, NASA Langley Research Center Hampton, VA 23665-5225		10. SPONSORING / MONITORING AGENCY REPORT NUMBER NASA CR-189660		
9. SPONSORING / MONITORING AGENCY NAME(S) AND ADDRESS(ES) National Aeronautics and Space Administration Langley Research Center Hampton, VA 23665-5225		11. SUPPLEMENTARY NOTES Langley Technical Monitor: Michael F. Card Final Report		
12a. DISTRIBUTION / AVAILABILITY STATEMENT Unclassified - Unlimited Subject Category 59		12b. DISTRIBUTION CODE		
13. ABSTRACT (Maximum 200 words) This report summarizes research conducted at the Institute for Computer Applications in Science and Engineering in applied mathematics, numerical analysis, and computer science during the period October 1, 1991 through March 31, 1992.				
14. SUBJECT TERMS Mathematics; Numerical Analysis; Computer Science			15. NUMBER OF PAGES 48	
			16. PRICE CODE A03	
17. SECURITY CLASSIFICATION OF REPORT Unclassified	18. SECURITY CLASSIFICATION OF THIS PAGE Unclassified	19. SECURITY CLASSIFICATION OF ABSTRACT	20. LIMITATION OF ABSTRACT	